

BY MARK HOLDAWAY • XIGNAL TECHNOLOGIES AG

Designing antialias filters for ADCs

CONTINUOUS-TIME ADCs CAN BENEFIT SIGNAL-CHAIN DESIGN. AN OVERVIEW OF DISCRETE- AND CONTINUOUS-TIME SYSTEMS DETAILS THE DIFFERENCES.

Nyquist-sampling theory lies at the heart of today's digital-communications systems. It requires that data-conversion systems include antialiasing input filters. Designers need to understand the requirements for antialiasing filters and examine the consequences of filter application. They must also consider the benefits of a new class of ADC that uses a low-power, high-speed, continuous-time-sampling method. These devices claim the ability to achieve a first Nyquist-zone-sampling capability without the aid of external filters.

You can reconstruct a time-continuous signal from discrete-time-sampled data if the original sampling rate is twice that of the highest frequency component in the sampled signal. The Nyquist-sampling theory states that data clocked with a sample rate of f_s (sampling frequency) samples/sec can effectively represent a signal of bandwidth as high as $0.5 \times f_s$ Hz. The Nyquist theory places demands on the sampling function, time, and amplitude precision. Sampling signals with signal content greater than a $0.5 \times f_s$ Hz bandwidth cause aliasing, a nonlinear process that results in frequency shifting. Signal content at frequencies greater than $0.5 \times f_s$ Hz folds around $0.5 \times f_s$ Hz—the Nyquist frequency—and alias back into the baseband. This aliasing creates a serious

problem: Once you sample the signal, you have no way of determining which resulting signal components originate from the desired signal band and which ones are aliased errors. Figure 1 shows two alias signals, A', a single tone, and B', a spectrum, each folding down into the first Nyquist zone. Note that A' originates in Nyquist Zone 4, and B' is from Zone 3. Also note that, in a communications application, this folding may allow interference signals to completely obscure information-bearing Signal A.

You should bandlimit a signal for digitization to eliminate any signal power beyond the frequency range of interest. The design of a suitable antialiasing-filter network may seem fairly trivial; however, as ADC linearity and performance improve, these filters become a significant part of the total system design.

IDEAL AND PRACTICAL FILTERS

Ideal baseband, lowpass antialiasing filters should have a steep transition band, excellent gain flatness, and low distortion in the passband—difficult goals to achieve. Furthermore, the stopband attenuation should be enough to reduce any residual out-of-band signal power to a level invisible to the ADC. You achieve this performance by employing stopband attenuation in excess of the dynamic range of the ADC (Figure 2). Assume that the stopband extends to infinity. Applications encountering high noise levels, especially those with high levels of interference occurring close to the edge of the first Nyquist zone, require filters with aggressive falloff. You achieve this performance using high-order filters that typically exhibit poor phase performance and result in dispersion or large group delay. In antialiasing filters, filtering takes place before the time-sampling point, or quantizer; these filters consequently require the use of an analog filter. This requirement is unfortunate because you can more easily and cost-effectively implement aggressive filters in the digital domain. High-order analog filters provide low harmonic distortion and gain flatness to in-band signals. However, the design of these filters is complex because they are too sensitive to gain matching to be practical at more than a few orders of attenuation magnitude. Furthermore, any passband harmonic distortion the filter introduces also produces undesirable

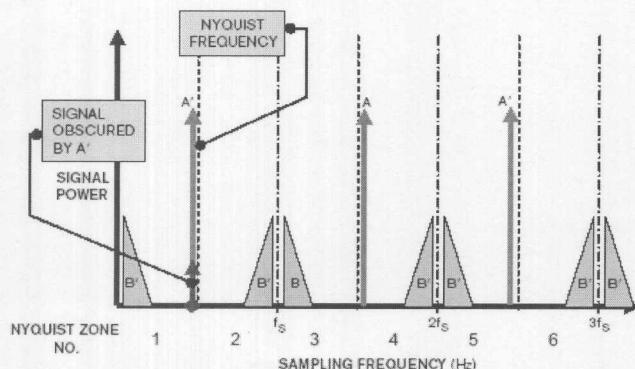


Figure 1 Alias signals A', a single tone, and B', a spectrum, can reside in any Nyquist zone if no antialias filter exists in a sampled system, but you can find both in Zone 1, where A' now obscures an information-bearing tone. A originates in Nyquist Zone 4, and B is from Zone 3.

signals in the output spectrum of the ADC. Insertion loss might also be important when using passive filters, which increase system noise.

An ideal antialiasing filter features 0-dB unity gain in the passband with little or no gain variation and a level of alias attenuation that matches the theoretical dynamic range of the data-conversion system in use. You derive a first approximation of this value from the theoretical SNR (signal-to-noise ratio) for an N-bit ADC: $\text{SNR} = 6.02 \times N + 1.76$ dB. For a 14-bit ADC, this approximation requires 80- to 86-dB attenuation with an ideal SNR of approximately 86 dB.

A number of standardized filter-transfer functions, including Bessel, Butterworth, Chebyshev, and elliptic, exist. Each has specific characteristics in the passband, transition band, and stopband. Selecting the appropriate topology depends on the most critical performance aspects of a design. Butterworth filters have the flattest passband region and minimal group delays. Chebyshev filters have steeper roll-offs but more passband ripple. Elliptic filters feature the steepest roll-off (Figure 3). The figure does not show a Bessel filter, which has a more gradual roll-off but has the key advantage of a linear, or constant, phase response. A number of public-domain tools exist to help developers in the design of a suitable antialiasing filter.

A consequence of using an antialiasing filter is the limit on available alias-free bandwidth when you use it in a traditional ADC. At first glance, the Nyquist theorem seems to promise a lot. Consider an ADC that samples at 40M samples/sec at a clock frequency of 40 MHz. It theoretically promises a 20-MHz signal bandwidth. However, aliasing with practical filter design means that the free bandwidth is considerably less than this amount. A 14-bit converter can resolve to one part in 2^14 —that is, one part in 16,384. To bury any alias component in the ADC's noise floor requires attenuation to be less than ± 0.5 LSB. That amount equates to 90-dB attenuation—that is, ± 0.5 LSB = one part in $32,768 = 90.3$ dB. In practical terms, however, this level of attenuation need exceed only the measured SNR of a 14-bit ADC. A more realistic level in the filter design is an attenuation of 80 dB.

Figure 3 shows several possible filter topologies, including two Butterworth-transfer functions—those of four- and eight-pole systems—both compared with an ideal Nyquist filter. Note that, by convention, the cutoff frequency is the point at which the filter produces 3 dB of attenuation. The horizontal axis shows the normalized input frequency as a ratio of the absolute frequency to the cutoff frequency. Note that the four-pole curve does not drop to 80 dB until the input frequency has risen to 10 times the cutoff frequency.

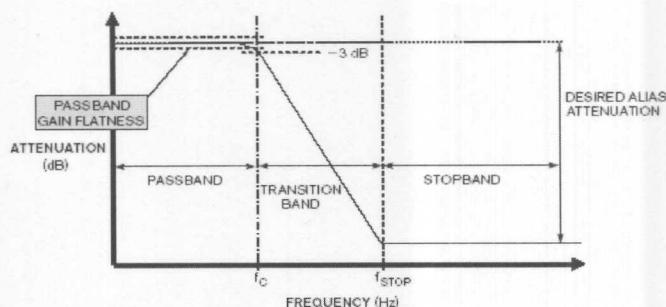


Figure 2 Ideal lowpass antialiasing filters should have a steep transition band, excellent gain flatness, and low distortion in the passband.

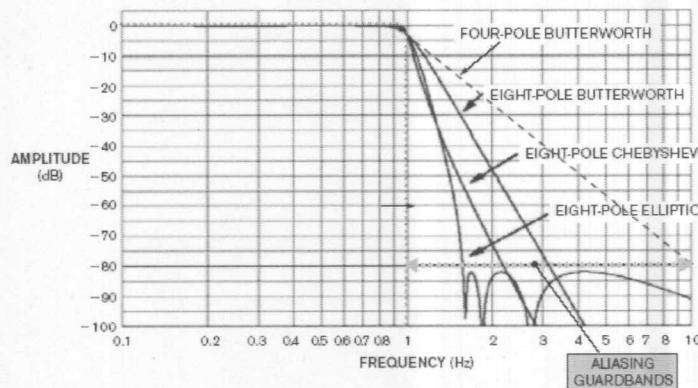


Figure 3 Possible antialias-filter designs illustrate the different transition-band characteristics of example filter systems.

So, if this ADC were to sample a 5-MHz signal, then this system would still see frequencies all the way out to 50 MHz. To fully sample the 5-MHz bandwidth and eliminate aliasing, the correct sample frequency using this filter would need to be 100 MHz. The range of 5 to 50 MHz becomes a guardband against alias errors. An obvious option is to look for higher performance filters. Consider an aggressive, eight-pole filter. Inspection shows that the 80-dB-attenuation point occurs at a frequency that is 3.2 times the cutoff frequency, or 16 MHz, a significantly reduced alias guardband. Alias-free sampling requires considerably more system bandwidth to handle the alias-guardband needs of an application. It is also important to note the cost trade-offs you must weigh when considering the severity of the antialiasing filter and the performance level of the ADC.

To ease antialiasing-filter design, pipeline ADCs—often confusingly referred to as Nyquist converters—have been offering increased sample rates and input bandwidths. Oversampling a signal at twice the Nyquist rate evenly spreads the ADC's quantization-noise power into a two-

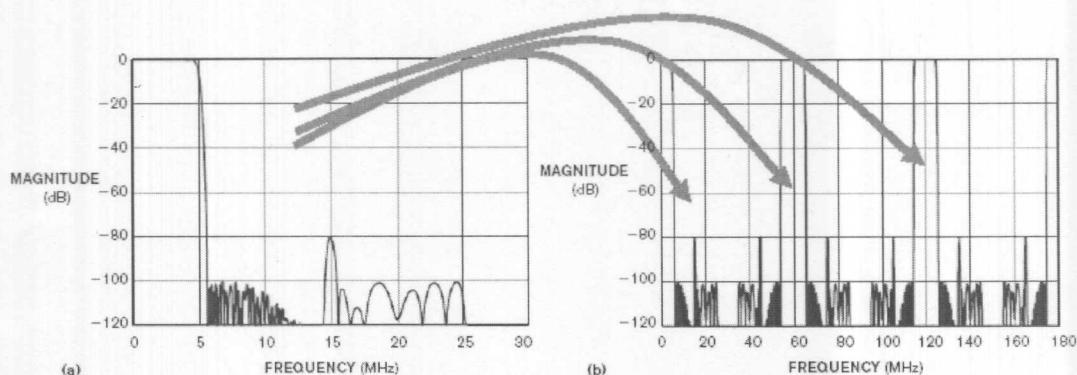


Figure 4 Discrete time sampling produces a lowpass signal-transfer function in a discrete time delta-sigma ADC (a). The graph in (a) appears to show alias protection; however, the transfer function of (a) is wrapped around integer multiples of the sampling frequency, as the expanded plot (b) shows. Aliasing gaps appear centered on 60, 120, and 180 MHz in this case.

times-wider frequency band. Applying decimation to sub-sample the resultant output samples yields a 3-dB/octave conversion gain. This technique is useful for deployment in delta-sigma converters because it not only produces dynamic-range improvements, but also reduces the pressure on the antialiasing filter by relaxing filter roll-off. Lower order antialiasing filters are easier to match across multiple channels than higher order ones. Oversampling techniques reduce the demands on the filter networks, but higher-sample-rate ADCs and faster digital processing use more power and increase cost.

You must also consider the phase response of the antialiasing filters. A filtered signal should not see any significant phase alteration. This alteration becomes even worse if phase varies according to input frequency. You normally measure phase variation in a filter in terms of group delay—that is, the derivative of phase with respect to frequency. For a nonconstant group delay, a signal spreads out in time, causing poor impulse response. Dispersion may be an additional worry for system performance. This factor is important in the design of

ultrasound systems in which the received-signal phase carries reflection information.

DELTA-SIGMA CONVERTERS

Delta-sigma techniques place lower demands on antialiasing filters. Delta-sigma converters exploit oversampling. In the past, designers improved dynamic range by using high oversampling rates and a simple low-resolution quantizer. However, simple oversampling produces minimal conversion-gain improvements. Applying feedback provides a faster route to conversion-gain improvements.

Delta-sigma modulators apply feedback to shape the quantization noise in the frequency domain by pushing most noise power into frequencies beyond the signal band of interest. Filtering can reduce the noise power in this band. Employing oversampled systems, which provide free frequency space beyond the signal band of interest, accomplishes this goal. Conventional Nyquist converters achieve a 3-dB/octave conversion gain through $2\times$ oversampling. Delta-sigma converters more efficiently build conversion gain, which the order of the applied feedback loop determines. First-, second-, or third-order loops can provide 9-, 15-, or 21-dB/octave conversion gain, respectively.

Most delta-sigma-converter implementations are discrete-time systems in which designers build the loop-filter components from simple switched-capacitor filters. The signal-transfer function of a delta-sigma modulator is an important factor in such a design. Signal-transfer performance looks promising in traditional discrete-time systems. Digital-decimation filters define the effective passband and provide a sharp transition band. Unfortunately, switched-capacitor-filter networks, which define the input bandwidth, add a discrete-sampling effect to the modulator structure. This discrete sampling causes a lowpass signal-transfer function (Figure 4a). Although this function seems acceptable, a closer inspection of a wideband-frequency plot reveals a problem: The passband of the digital filter wraps around integer multiples of the sample frequency at 60, 120, and 180 MHz (Figure 4b). No alias attenuation whatsoever exists at these points, and this characteristic extends to infinity. Preventing high-level, out-of-band noise at multiples

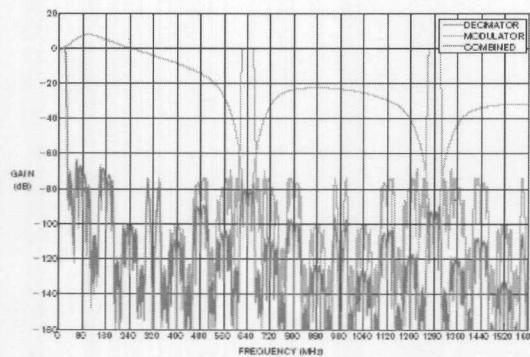


Figure 5 An aliasing-mitigation system ensures the analog-loop filter provides maximum stopband attenuation at the oversampling frequency of the modulator.

of the oversample rate is a challenge and a downside of such designs.

CONTINUOUS TIME

In a continuous-time modulator, you implement noise shaping using conventional analog active filters. The benefit of the continuous-time approach is that you can design the loop filter to handle alias filtering of the input signal. Tai-

loring this filter system for a specific product, the maximum-loop-filter attenuation coincides with the minimum attenuation that the decimation filter offers. An aliasing-mitigation system ensures that the analog-loop filter provides maximum stopband attenuation at the oversampling frequency of the modulator (Figure 5, green line). This attenuation ensures that no noise power

beyond the oversampling frequency can enter the first Nyquist zone. The back-end digital filter provides a sharp stopband attenuation, limiting the maximum effective input bandwidth of the ADC (blue line). Through this arrangement, the maximum analog-loop-filter attenuation always coincides with the folded-digital-filter minimum to maintain a high level of wideband attenuation. The maximum attenuation of the analog-loop filter coincides with the alias passband of the digital filter. The purple line shows the composite transfer function.

The specific implementation of a given delta-sigma topology determines the performance of the antialias system. For example, the 14-bit-resolution, 20M- to 40M-sample/sec Xignal (www.xignal.com) XT11400 ADC achieves 76-dB SNR and provides a 20-MHz analog-input bandwidth. The passband gain flatness is ± 0.002 dB, the transition band is approximately 2.5 MHz wide, and the unit achieves alias attenuation of 80 dB beyond 22.5 MHz, all without any external filtering. A digital allpass-filter stage, which reduces dispersion to 0.3 samples, minimizes group delay. Such approaches have benefits in reducing design complexity, especially in multiple-channel designs in which cross-channel filter matching is a major issue.

In summary, delta-sigma modulators use oversampling to help simplify anti-aliasing-filter design. For discrete-time systems, you must use caution in designing anti-aliasing filters because of the potential occurrence of high-frequency noise, which can couple and fold directly into the baseband. A continuous-time alternative can eliminate the need for all external anti-aliasing filters. The maximum attenuation of the analog-loop filter aims successfully to intercept the passband frequency of the digital-decimation filter at the oversampling frequency. EDN

AUTHOR'S BIOGRAPHY

Mark Holdaway is product-marketing director at Xignal Technologies AG (Unterhaching, Germany). He holds a bachelor's degree in electronics engineering from the University of Salford (England). His interests include cycling, skiing, gardening, reading, and writing. You can reach him at mark.holdaway@xignal.com.

WORKING HARD?

Little Things Make the Hard Work Easier

Count on Keystone for Interconnect Components & Electronic Hardware

Our Catalog also includes:

- Battery Clips, Contacts & Holders
- Fuse Clips & Holders • Spacers & Standoffs
- Terminals & Test Points • Panel Hardware
- PC Board Hardware • Terminal Boards
- Pins, Plugs, Jacks & Sockets
- Multi-Purpose Hardware

Available from our global distributor network

RoHS Compliant - ISO 9001 Certified

KEYSTONE
ELECTRONICS CORP.

(718) 956-8900 • (800) 221-5510 • FAX (718) 956-9040
e-mail: kec@keyelco.com • Website: keyelco.com

Multi-Purpose Hardware